CS543 Distributed Systems
QE, 2015

1. With the advancement of network and computing technologies, users can ubiquitously store and retrieve their data without relying on a specific server. However, they are still locked in a certain vendor service. To overcome this, there have been several research efforts to build a distributed storage service where storage servers are replicated across multiple places, support users at proximity, and propagate updates to other replicas.
   a) It is said that consistency and performance cannot stand together in this design. Discuss why.
   b) Discuss what kind of a naming scheme should be designed to support the target service (i.e. a user's data (i.e. file) can be accessed regardless of its current location)
   c) Discuss what kind of a data consistency protocol can implement the target service and why the protocol provides no guarantee of sequential consistency and fault tolerance if implemented as a non-blocking.